Chapter 15
Configure Ethernet Interfaces

Ethernet was developed in the early 1970s at the Xerox Palo Alto Research Center as a data-link control layer protocol for interconnecting computers. It was first widely used at 10 Mbps over coaxial cables and later over unshielded twisted pairs using 10BaseT. More recently, 100BaseTX (Fast Ethernet, 100 Mbps) and Gigabit Ethernet (1 Gbps) have become available.

Juniper Networks routers support the following types of Ethernet interfaces:

- Fast Ethernet
- Gigabit Ethernet
- Management Ethernet interface, which is an out-of-band management interface within the router
- Internal Ethernet interface, which connects the Routing Engine to the Packet Forwarding Engine
- Aggregated Ethernet interface, a logical linkage of Fast Ethernet or Gigabit Ethernet physical connections

This chapter discusses the following topics specific to configuring the different types of Ethernet interfaces in the router:

Configure Ethernet Physical Interface Properties

To configure Fast Ethernet-specific physical interface properties, include the fastether-options statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
link-mode (full-duplex | half-duplex);
speed (10m | 100m);
vlan-tagging;
fastether-options {
    802.3ad aex;
    (flow-control | no-flow-control);
    ingress-rate-limit rate;
    (loopback | no-loopback);
    source-address-filter {
        mac-address;
    }
    (source-filtering | no-source-filtering);
}
```

The statement speed (10m | 100m) applies only to the management Ethernet interface (fxp0) and to the Fast Ethernet 12-port and 48-port PICs. The 4-port and 8-port Fast Ethernet PICs support a speed of 100 Mbps only.

To configure Gigabit Ethernet-specific physical interface properties, include the gigether-options statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
gigether-options {
    802.3ad aex;
    (flow-control | no-flow-control);
    (loopback | no-loopback);
    source-address-filter {
        mac-address;
    }
    (source-filtering | no-source-filtering);
}
```

To configure aggregated Ethernet-specific physical interface properties, include the aggregated-ether-options statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
aggregated-ether-options {
    (flow-control | no-flow-control);
    link-speed speed;
    (loopback | no-loopback);
    minimum-links number;
    source-address-filter {
        mac-address;
    }
    (source-filtering | no-source-filtering);
}
```

You can configure the following properties specific to aggregated Ethernet, Fast Ethernet, or Gigabit Ethernet interfaces:

Configure Ethernet Link Aggregation

On Fast Ethernet and Gigabit Ethernet interfaces, you can associate a physical interface with an aggregated Ethernet interface. To enable the aggregated link, include the 802.3ad statement at the [edit interfaces *interface-name* fastether-options] or [edit interfaces *interface-name* gigether-options] hierarchy level:

```
802.3ad aex;
```

You specify the interface instance number *x* to complete the link association; *x* can range from 0 through 15, for a total of 16 aggregated interfaces. You must also include a statement defining ae*x* at the [edit interfaces] hierarchy level. For more information, see “Configure Aggregated Ethernet Interfaces” on page 247. You can optionally specify other physical properties that apply specifically to the aggregated Ethernet interfaces; for details, see “Configure Ethernet Physical Interface Properties” on page 230, and for a sample configuration, see “Example: Configure Aggregated Ethernet Interfaces” on page 250.

Configure Aggregated Ethernet Link Speed

On aggregated Ethernet interfaces, you can set the required link speed for all interfaces included in the bundle. All interfaces that make up a bundle must be the same speed. If you include in the aggregated Ethernet interface an individual link that has a speed different from the speed you specify in the link-speed parameter, an error message will be logged. To set the required link speed, include the link-speed statement at the [edit interfaces *interface-name* aggregated-ether-options] hierarchy level:

```
[edit interfaces interface-name aggregated-ether-options]
link-speed speed;
```

speed can be in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

Configure Aggregated Ethernet Minimum Links

On aggregated Ethernet interfaces, you can configure the minimum number of links that must be up for the bundle as a whole to be labeled up. To configure the minimum number, include the minimum-links statement at the [edit interfaces *interface-name* aggregated-ether-options] hierarchy level:

```
[edit interfaces interface-name aggregated-ether-options]
minimum-links number;
```

By default, minimum-links has a value of 1. *number* can be a value from 1 through 8.

Configure MAC Address Filtering

On aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces, you can enable source address filtering, which blocks all incoming packets to that interface. To enable the filtering, include the source-filtering statement at the [edit interfaces *interface-name* aggregated-ether-options], [edit interfaces *interface-name* fastether-options], or [edit interfaces *interface-name* gigether-options] hierarchy level:

```
source-filtering;
```

When source address filtering is enabled, you can configure the interface to receive packets from specific MAC addresses. To do this, specify the MAC addresses in the source-address-filter statement at the [edit interfaces *interface-name* aggregated-ether-options], [edit interfaces *interface-name* fastether-options], or [edit interfaces *interface-name* gigether-options] hierarchy level:

```
source-address-filter {
   mac-address;
   <additional-mac-address;>
}
```

You can specify the MAC address as *nn:nn:nn:nn:nn:nn* or *nnnn.nnnn.nnnn*, where *n* is a hexadecimal number. To specify more than one address, include multiple *mac-address* statements in the source-address-filter statement.

If the remote Ethernet card is changed, the interface will not be able to receive packets from the new card because it will have a different MAC address.

Support for source address filters is limited on the Fast Ethernet 12-port and 48-port PIC interfaces.

Configure Loopback Mode

By default, local aggregated Ethernet, Fast Ethernet, or Gigabit Ethernet interfaces connect to a remote system. To place an interface in loopback mode, include the loopback statement at the [edit interfaces *interface-name* aggregated-ether-options], [edit interfaces *interface-name* fastether-options], or [edit interfaces *interface-name* gigether-options] hierarchy level:

```
loopback;
```

To return to the default—that is, to disable loopback mode—delete the loopback statement from the configuration:

```
[edit]
user@host# delete interfaces fe-fpc/pic/port fastether-options loopback
```

To explicitly disable loopback mode, include the no-loopback statement at the [edit interfaces *interface-name* aggregated-ether-options], [edit interfaces *interface-name* fastether-options], or [edit interfaces *interface-name* gigether-options] hierarchy level:

```
no-loopback;
```

Configure Flow Control

By default, the router imposes flow control to regulate the amount of traffic sent out a Fast Ethernet or Gigabit Ethernet interface. This is useful if the remote side of the connection is a Fast Ethernet or Gigabit Ethernet switch.

You can disable flow control if you want the router to permit unrestricted traffic. To disable flow control, include the no-flow-control statement at the [edit interfaces *interface-name* aggregated-ether-options], [edit interfaces *interface-name* fastether-options], or [edit interfaces *interface-name* gigether-options] hierarchy level:

```
no-flow-control;
```

To explicitly reinstate flow control, include the flow-control statement at the [edit interfaces *interface-name* aggregated-ether-options], [edit interfaces *interface-name* fastether-options], or [edit interfaces *interface-name* gigether-options] hierarchy level:

```
flow-control;
```

Configure the Link Characteristics

By default, the router's management Ethernet interface, fxp0, autonegotiates whether to operate in full-duplex or half-duplex mode. Fast Ethernet interfaces can operate in either full-duplex or half-duplex mode, and all other interfaces can operate only in full-duplex mode. For Gigabit Ethernet, the link partner must also be set to full duplex.

To explicitly configure an Ethernet interface to operate in either full-duplex or half-duplex mode, include the link-mode statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
link-mode (full-duplex | half-duplex);
```

Configure the Interface Speed

On Fast Ethernet 12-port and 48-port PIC interfaces and the management Ethernet interface (fxp0) only, you can explicitly set the interface speed to either 10 Mbps or 100 Mbps.

To explicitly configure the speed on an interface, include the speed statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
speed (10m | 100m);
```

Configure the Ingress Rate Limit

On Fast Ethernet 8-port, 12-port, and 48-port PIC interfaces only, you can apply port-based rate limiting to the ingress traffic that arrives at the PIC.

To configure an ingress rate limit on a Fast Ethernet 8-port, 12-port, or 48-port PIC interface, include the ingress-rate-limit statement at the [edit interfaces *interface-name* fastether-options] hierarchy level:

```
[edit interfaces interface-name fastether-options]
ingress-rate-limit rate;
```

rate can range in value from 1 through 100 Mbps.

Configure 802.1Q VLANs

For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces, the software supports a subset of the IEEE 802.1Q standard for channelizing an Ethernet interface into multiple logical interfaces, allowing many hosts to be connected to the same Gigabit Ethernet switch, but preventing them from being in the same routing or broadcast domain.

The software supports receiving and forwarding routed Ethernet frames with 802.1Q Virtual Local Area Network (VLAN) tags, and running VRRP over 802.1Q-tagged interfaces. To configure the router to receive and forward frames with 802.1Q VLAN tags, include the vlan-tagging statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
vlan-tagging;
```

Gigabit Ethernet interfaces can be partitioned; you can assign up to 4095 different logical interfaces, one for each VLAN, but you are limited to a maximum of 1024 VLANs on any single Gigabit Ethernet port. Fast Ethernet interfaces can also be partitioned, with a maximum of 1024 logical interfaces for the four-port FE PIC and 16 logical interfaces for the M40e and M160 FE-48 PIC. Table 17 lists VLAN ID range by interface type.

Table 17: VLAN ID Range by Interface Type

Interface Type	VLAN ID Range
4-port, 8-port, and 12-port Fast Ethernet	0 through 1023
48-port Fast Ethernet	0 through 4094
Gigabit Ethernet	0 through 4094
Management and internal Ethernet interfaces	0 through 1023

To bind a VLAN ID to a logical interface, include the vlan-id statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
vlan-id number;
```

Because IS-IS has an 8-bit limit for broadcast multiaccess media, you cannot set up more than 255 adjacencies over Gigabit Ethernet using VLAN tagging. For further information on IS-IS capabilities, see the *JUNOS Internet Software Configur ation Guide: R outing and R outing Protocols*.

Configure VLAN CCC Encapsulation

Ethernet interfaces with VLAN tagging enabled can use VLAN circuit cross-connect (CCC) encapsulation. To configure the encapsulation on a physical interface, include the encapsulation vlan-ccc statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
encapsulation vlan-ccc;
```

Ethernet interfaces in VLAN mode can have multiple logical interfaces, but in CCC mode VLAN IDs from 0 through 511 are reserved for normal VLANs, and VLAN IDs from 512 through 1023 are reserved for CCC VLANs.

In general, you configure an interface's encapsulation at the [edit interfaces *interface-name*] hierarchy level. However, for some encapsulation types, including Ethernet VLAN CCC, you also can configure the encapsulation type that is used inside the VLAN circuit itself. To do this, include the encapsulation statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
encapsulation vlan-ccc;
```

You cannot configure a logical interface with an encapsulation of VLAN CCC unless you also configure the physical device with the same encapsulation. The logical interface must also have a VLAN ID in the range from 512 through 4094; if the VLAN ID is 511 or lower, it will be subject to the normal destination filter lookups in addition to source address filtering.

Example: Configure VLAN CCC Encapsulation

Configure VLAN CCC encapsulation on a Gigabit Ethernet interface:

```
interfaces ge-2/1/0 {
    vlan-tagging;
    encapsulation vlan-ccc;
    unit 0 {
        encapsulation vlan-ccc;
        vlan-id 600;
    }
}
```

Configure Extended VLAN Cross-Connect Encapsulation

Gigabit Ethernet interfaces with VLAN tagging enabled can use extended VLAN circuit cross-connect (CCC), which allows 802.1Q tagging or translational cross-connect (TCC) encapsulation, which allows circuits to have different media on either side of the connection. To configure the encapsulation on a physical interface, include the encapsulation extended-vlan-ccc or encapsulation extended-vlan-tcc statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
encapsulation (extended-vlan-ccc | extended-vlan-tcc);
```

One-port Gigabit Ethernet, two-port Gigabit Ethernet, four-port Gigabit Ethernet, and four-port Fast Ethernet PICs can use Extended Ethernet TCC encapsulation. Extended Ethernet TCC is not supported on the T-series platforms.

For extended VLAN CCC encapsulation, all VLAN IDs from 0 through 4094 are valid. For extended VLAN TCC encapsulation, all VLAN IDs from 0 through 1024 are valid.

Extended VLAN CCC is not supported on four-port Gigabit Ethernet PICs.

For extended VLAN CCC, the VLAN IDs on ingress and egress interfaces must be the same. For back-to-back connections, all VLAN IDs must be the same.

Example 1: Configure Extended VLAN CCC Encapsulation

Configure extended VLAN CCC encapsulation on Gigabit Ethernet ingress and egress interfaces:

```
interfaces ge-0/0/0 {
    vlan-tagging;
    encapsulation extended-vlan-ccc;
    unit 0 {
        vlan-id 2;
        family ccc;
    }
}

interfaces ge-1/0/0 {
    vlan-tagging;
    encapsulation extended-vlan-ccc;
    unit 0 {
        vlan-id 2;
        family ccc;
    }
}
```

Example 2: Configure Extended VLAN TCC Encapsulation

Configure a full-duplex Layer 2.5 translational cross-connect between Router A and Router C, using a Juniper Networks router, Router B, as the TCC interface. Extended VLAN TCC encapsulation provides an Ethernet wide area circuit for interconnecting IP traffic. See the topology in Figure 16.

Figure 16: Example Topology of Layer 2.5 Translational Cross-Connect

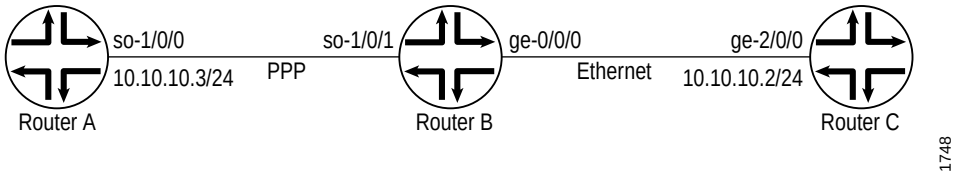


```
interfaces ge-0/0/0 {
    vlan-tagging;
    encapsulation extended-vlan-tcc;
    unit 0 {
        vlan-id 1;
        family tcc {
            remote {
                mac-address 0011.2233.4455;
            }
        }
    }
}
```

The Router A–to–Router B circuit is PPP, and the Router B–to–Router C circuit is Ethernet with VLAN tagging enabled.

Configure Static ARP Table Entry for Router C

Extended VLAN TCC does not look at Layer 3 IP addresses; therefore, for the above example to work, you must configure a static ARP table entry, defining a mapping between the IP and MAC addresses of Router C, as shown in the following example:

```
interfaces ge-2/0/0 {
    unit 0 {
        family inet {
            address 10.10.10.2/24; {
                arp 10.10.10.3 mac 0011.2233.4455;
            }
        }
    }
}
```

0011.2233.4455 is the MAC address of Router B's ge-0/0/0 interface.

For more information, see "Configure Ethernet TCC and Extended VLAN TCC" on page 143.

For more information about static ARP, see "Configure Static ARP Table Entries" on page 238.
For more information about Ethernet TCC, see the *JUNOS Internet Softw are Configur ation Guide: VPNs* .

Configure Static ARP Table Entries

For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces, you can configure static ARP table entries, defining mappings between IP and MAC addresses. To configure static ARP table entries, include the arp statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address]
arp ip-address (mac | multicast-mac) mac-address <publish>;
```

The IP address that you specify must be part of the subnet defined in the enclosing address statement.

To associate a multicast MAC address with a unicast IP address, include the multicast-mac statement.

Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn.nnnn.nnnn* or *nn:nn:nn:nn:nn:nn*. For example, 0011.2233.4455 or 00:11:22:33:44:55.

For unicast MAC addresses only, if you include the publish option, the router replies to proxy ARP requests.

Example: Configure Static ARP Table Entries

Configure two static ARP table entries on the router's management interface:

```
interfaces fxp0 {
    unit 0 {
        family inet {
            address 10.10.0.11/24 {
                arp 10.10.0.99 mac 0001.0002.0003;
                arp 10.10.0.101 mac 00:11:22:33:44:55 publish;
            }
        }
    }
}
```

Configure VRRP

For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces, you can configure the Virtual Router Redundancy Protocol (VRRP). VRRP allows hosts on a LAN to make use of redundant routers on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routers share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routers is the master (active) and the others are backups. If the master fails, one of the backup routers becomes the new master, thus always providing a virtual default router and allowing traffic on the LAN to be routed without relying on a single router.

VRRP is defined in the following document:

RFC 2338, *Virtual Router Redundanc y Protocol*

To configure VRRP, include the vrrp-group statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address]
vrrp-group group-number {
    virtual-address [ addresses ];
    priority number;
    (accept-data | no-accept-data);
    advertise-interval seconds;
    authentication-type authentication;
    authentication-key key;
    (preempt | no-preempt);
    track {
        interface interface-name priority-cost cost;
    }
}
```

To trace VRRP operations, include the traceoptions statement at the [edit protocols vrrp] hierarchy level:

```
[edit protocols vrrp traceoptions]
    file {
        filename filename;
        files number;
        size size;
        (world-readable | no-world-readable);
    }
    flag flag;
```

For more information, see “Trace VRRP Operations” on page 243.

You can configure the following VRRP properties:

Configure Basic VRRP Support on page 240

Configure VRRP Authentication on page 241

Configure the Advertisement Interval for the VRRP Master Router on page 241

Configure a Backup Router to Preempt the Master Router on page 242

Accept Packets Destined for the Virtual IP Address on page 242

Configure an Interface to Be Tracked on page 243

Trace VRRP Operations on page 243

For a VRRP configuration example, see “Example: Configure VRRP” on page 244.

Configure Basic VRRP Support

To configure basic VRRP support, configure VRRP groups on interfaces by including the following statements at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address]
vrrp-group group-number {
    virtual-address [ addresses ];
    priority number;
}
```

An interface can be a member of one or more VRRP groups. For each group, you must configure the following:

Group number—Identifies the VRRP group. It can be a value from 0 through 255.

If you also enable MAC source address filtering on the interface, as described in "Configure MAC Address Filtering" on page 79, you must include the virtual MAC address in the list of source MAC addresses that you specify in the source-address-filter statement. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

Addresses of one or more virtual routers that are members of the VRRP group—Virtual IP addresses associated with the virtual router in the VRRP group. Normally, you configure only one virtual IP address per group. The virtual IP addresses must be the same for all routers in the VRRP group.

In the addresses, specify the address only. Do not include a prefix length.

If you configure a virtual IP address to be the same as the interface's address (the address configured with the address statement), the interface becomes the master virtual router for the group. In this case, you must configure the priority to be 255 and you must configure preemption by including the preempt statement. If you have multiple VRRP groups on an interface, the interface can be the master virtual router for only one of the groups.

If the virtual IP address you choose is not the same as the interface's address, you must ensure that this address does not appear anywhere else in the router's configuration. Check that you do not use this address for other interfaces, for the IP address of a tunnel, or for the IP address of static ARP entries.

Priority for this router to become the master virtual router—Value used to elect the master virtual router in the VRRP group. It can be a number from 1 through 255. The default value for backup routers is 100. A larger value indicates a higher priority. The router with the highest priority within the group becomes the master router.

Within a single VRRP group, the master and backup routers cannot be the same router.

Configure VRRP Authentication

All VRRP protocol exchanges can be authenticated to guarantee that only trusted routers participate in the AS's routing. By default, VRRP authentication is disabled. You can configure one of the following authentication methods; each VRRP group must use the same method:

Simple authentication—Uses a text password included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet.

MD5 algorithm—Creates the authentication data field in the IP authentication header. This header is used to encapsulate the VRRP protocol data unit (PDU). The receiving router uses an authentication key (password) to verify the authenticity of the IP authentication header and VRRP PDU.

To enable authentication and specify an authentication method, include the authentication-type statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address
vrrp-group group-number]
authentication-type authentication;
```

authentication can be none, simple, or md5. The authentication type must be the same for all routers in the VRRP group.

If you included the authentication-type statement to select an authentication method, you can configure a key (password) on each interface by including the authentication-key statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address
vrrp-group group-number]
authentication-key key;
```

The key (password) is an ASCII string. For simple authentication, it can be 1 through 8 characters long. For MD-5 authentication, it can be 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (“ “). The key must be the same for all routers in the VRRP group.

Configure the Advertisement Interval for the VRRP Master Router

By default, the master router sends VRRP advertisement packets every second to all members of the VRRP group. These packets indicate that the master router is still operational. If the master router fails or becomes unreachable, the backup router with the highest priority value becomes the new master router.

To modify the time between the sending of VRRP advertisement packets, include the advertise-interval statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address
vrrp-group group-number]
advertise-interval seconds;
```

The interval can range from 1 through 255 seconds. The interval must be the same for all routers in the VRRP group.

Configure a Backup Router to Preempt the Master Router

By default, a higher priority backup router preempts a lower priority master router. To explicitly allow the master router to be preempted, include the preempt statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address
vrrp-group group-number]
preempt;
```

To prohibit a higher priority backup router from preempting a lower priority master router, include the no-preempt statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address
vrrp-group group-number]
no-preempt;
```

The router that owns the IP address(es) associated with the virtual router always preempts, independent of the setting of this flag.

Accept Packets Destined for the Virtual IP Address

To configure an interface to accept packets destined for the virtual IP address, include the accept-data statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address
vrrp-group group-number]
accept-data;
```

To prohibit the interface from accepting packets destined for the virtual IP address, include the no-accept-data statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address
vrrp-group group-number]
no-accept-data;
```

The accept-data statement has the following consequences:

You do not need to include the accept-data statement to activate this feature if the master router owns the virtual IP address.

If you do not include the accept-data statement, and if the master router owns the virtual IP address, the master router responds to ICMP message requests only.

You cannot include the accept-data statement when the priority of the master router is set to 255.

To restrict incoming IP packets to ICMP only, you must configure firewall filters to accept only ICMP packets.

If you include the accept-data statement, your router configuration will not comply with RFC 2338.

If you include the accept-data statement, VRRP clients should be able to process Gratuitous ARP.

If you include the accept-data statement, VRRP clients should not use packets other than ARP replies to update their ARP cache.

Configure an Interface to Be Tracked

VRRP can track whether an interface is up or down and dynamically change the priority of the VRRP group based on the state of the tracked interface, which might trigger a new master router election.

When interface tracking is enabled, you cannot configure a priority of 255, thereby designating the master router. For each VRRP group, 1 through 10 interfaces can be tracked.

To configure an interface to be tracked, include the track statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group
group-number]
track {
    interface interface-name priority-cost cost;
}
```

The priority cost is the value to be subtracted from the configured VRRP priority when the tracked interface is down, forcing a new master router election. The cost can range from 1 through 254. The sum of the costs for all tracked interfaces or routes must be less than or equal to the configured priority of the VRRP group.

Trace VRRP Operations

To trace VRRP operations, include the traceoptions statement at the [edit protocols vrrp] hierarchy level.

By default, VRRP logs the error, DCD configuration, and routing socket events in a file in the /var/log directory. By default, this file is named /var/log/vrrpd. The default file size is 1MB, and three files are created before the first one gets overwritten.

To change the configuration of the logging file, include the file statement at the [edit protocols vrrp traceoptions] hierarchy level:

```
[edit protocols vrrp traceoptions]
    file {
        filename filename;
        files number;
        size size;
        (world-readable | no-world-readable);
    }
    flag flag;
```

You can specify the following VRRP tracing flags:

all—Trace all VRRP operations.

database—Trace all database changes.

general—Trace all general events.

interfaces—Trace all interface changes.

normal—Trace all normal events.

packets—Trace all packets sent and received.

state—Trace all state transitions.

timer—Trace all timer events.

Example: Configure VRRP

Configure one master (Router A) and one backup (Router B) router. Note that the address configured in the virtual-address statements differs from the addresses configured in the address statements.

Router A:

```
[edit]
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 192.168.1.20/24 {
                    vrrp-group 27 {
                        virtual-address 192.168.1.15;
                        priority 254;
                        authentication-type simple;
                        authentication-key booJUM;
                    }
                }
            }
        }
    }
}
```

Router B:

```
[edit]
interfaces {
    ge-4/2/0 {
        unit 0 {
            family inet {
                address 192.168.1.24/24 {
                    vrrp-group 27 {
                        virtual-address 192.168.1.15;
                        priority 200;
                        authentication-type simple;
                        authentication-key booJUM;
                    }
                }
            }
        }
    }
}
```

When configuring multiple VRRP groups on an interface, configure one to be the master virtual router for that group:

```
[edit]
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 192.168.1.20/24 {
                    vrrp-group 2 {
                        virtual-address 192.168.1.20;
                        priority 255;
                        advertise-interval 3;
                        preempt;
                    }
                    vrrp-group 10 {
                        virtual-address 192.168.1.55;
                        priority 201;
                        advertise-interval 3;
                    }
                    vrrp-group 1 {
                        virtual-address 192.168.1.54;
                        priority 22;
                        advertise-interval 4;
                    }
                }
            }
        }
    }
}
```

Configure VRRP and MAC source address filtering on a Gigabit Ethernet interface. The VRRP group number is the decimal equivalent of the last byte of the virtual MAC address.

```
[edit interfaces]
ge-5/2/0 {
    gigether-options {
        source-filtering;
        source-address-filter {
            00:00:5e:00:01:0a;<— Virtual MAC address
        }
    }
    unit 0 {
        family inet {
            address 192.168.1.10/24 {
                vrrp-group 10 { <— VRRP group number
                    virtual-address 192.168.1.10;
                    priority 255;
                    preempt;
                }
            }
        }
    }
}
```

Configure the Management Ethernet Interface

The router's management Ethernet interface, fxp0, is an out-of-band management interface. You must configure an IP address and prefix length for this interface, which you commonly do when you first install the software:

```
[edit]
user@host# set interfaces fxp0 unit 0 family inet address/prefix-length
[edit]
user@host# show
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address/prefix-length;
            }
        }
    }
}
```

The management Ethernet interface must be configured for the router to function.

Configure the Internal Ethernet Interface

The internal Ethernet interface, fxp1, connects the Routing Engine with the System Control Board (SCB), System and Switch Board (SSB), Forwarding Engine Board (FEB), or Switching and Forwarding Module (SFM), depending on router model, in the Packet Forwarding Engine. The router software automatically configures this interface.

Do not modify or remove the configuration for the internal Ethernet interface that the software automatically configures. If you do, the router will stop functioning.

```
user@host> show configuration
...
interfaces {
...
   fxp1 {
      unit 0 {
         family tnp {
            address 1;
         }
      }
   }
}
```

Configure Aggregated Ethernet Interfaces

Link aggregation of Ethernet interfaces is defined in the IEEE 802.3ad standard. The JUNOS implementation of 802.3AD balances traffic across the member links within an aggregated Ethernet bundle based on the Layer 3 information carried in the packet. This implementation uses the same load balancing algorithm used for per-packet load balancing. For information about per-packet load balancing, see the *JUNOS Internet Softw are Guide: Routing and R outing Protocols*.

The JUNOS software does not provide load balancing for multicast traffic on aggregated interfaces. If a link carrying multicast data goes down, another link carries the traffic. This provides redundancy, not more bandwidth.

The JUNOS software does not support the Link Aggregation Control Protocol (LACP).

You configure an aggregated Ethernet virtual link by specifying the link number as a physical device and then associating a set of ports that have the same speed and are in full-duplex mode. The physical interfaces can be either Fast Ethernet or Gigabit Ethernet devices, but must not intermix within the same aggregated link.

To specify aggregated Ethernet interfaces, include the vlan-tagging statement at the [edit interfaces ae*x*] hierarchy level and include the vlan-id statement at the [edit interfaces ae*x* unit *logical-unit-number*] hierarchy level, as in the following example:

```
[edit interfaces]
ae0 {
    vlan-tagging;
    unit 0 {
        vlan-id 100;
        family inet {
            address 10.1.1.1/24;
        }
    }
}
```

By default, no aggregated Ethernet interfaces are created. You must define the number of aggregated Ethernet interfaces by including the device-count statement at the [edit chassis aggregated-devices ethernet] hierarchy level:

```
[edit chassis]
aggregated-devices {
    ethernet {
        device-count number;
    }
}
```

The maximum number of aggregated interfaces is 16, and the assigned number can range from 0 through 15. For information about configuring aggregated devices, see the *JUNOS Internet Software Configuration Guide: Getting Started*.

You must also specify the constituent physical links by including the 802.3ad statement at the [edit interfaces *interface-name* fastether-options] or [edit interfaces *interface-name* gigether-options] hierarchy level; for more information, see "Configure Ethernet Link Aggregation" on page 231. You can optionally specify other physical properties that apply specifically to the aggregated Ethernet interfaces; for details, see "Configure Ethernet Physical Interface Properties" on page 77. For a sample configuration, see "Example: Configure Aggregated Ethernet Interfaces" on page 250.

To delete an aggregated Ethernet interface from the configuration, issue the delete interfaces ae*x* command at the [edit] hierarchy level in configuration mode:

```
[edit]
user@host# delete interfaces aex
```

If you delete an aggregated Ethernet interface from the configuration, the software removes the configuration statements related to ae*x* and sets this interface to down state. However, the aggregated Ethernet interface is not deleted until you delete the chassis aggregated-devices ethernet device-count configuration statement.

Example: Configure Fast Ethernet Interfaces

The following configuration is sufficient to get a Fast Ethernet interface up and running. By default, IPv4 Fast Ethernet interfaces use 802.3 encapsulation.

```
[edit]
user@host# set interfaces fe-fpc/pic/port unit 0 family inet address local-address
user@host# show
interfaces {
    fe-fpc/pic/port {
        unit 0 {
            family inet {
                address local-address;
            }
        }
    }
}
```

Example: Configure Gigabit Ethernet Interfaces

The following configuration is sufficient to get a Gigabit Ethernet interface up and running. By default, IPv4 Gigabit Ethernet interfaces use 802.3 encapsulation.

```
[edit]
user@host# set interfaces ge-fpc/pic/port unit 0 family inet address local-address
user@host# show
interfaces {
    ge-fpc/pic/port {
        unit 0 {
            family inet {
                address local-address;
            }
        }
    }
}
```

The M160, T320, and T640 two-port Gigabit Ethernet PIC supports two independent Gigabit Ethernet links. This PIC is supported on the M160, T320, and T640 platforms only and it requires a Type 2 M160, Type 2 T320, or Type 2 T640 FPC.

Each of the two interfaces on the PIC is named:

```
ge-fpc/pic/[0.1]
```

Each of these interfaces has functionality identical to the Gigabit Ethernet interface supported on the single-port PIC.

Example: Configure Aggregated Ethernet Interfaces

The following set of configurations is sufficient to get an aggregated Gigabit Ethernet interface up and running.

```
[edit interfaces]
ae0 {
    vlan-tagging;
    unit 0 {
        vlan-id 100;
        family inet {
            address 10.1.1.1/24;
        }
    }
}

[edit chassis]
aggregated-devices {
    ethernet {
        device-count 15;
    }
}

[edit interfaces]
ge-1/3/0 {
    gigether-options {
        802.3ad ae0;
    }
}

[edit interfaces ae0]
aggregated-ether-options {
    link-speed 1g;
    minimum-links 5;
}
```